

Long-Term Evaluation of GPS Timing Receiver Failures

DIETER HÖCHTL, ULRICH SCHMID

Technische Universität Wien

Department of Automation

Treitlstraße 1, A-1040 Vienna, Austria

Email: {hoechtl, s}@auto.tuwien.ac.at

Abstract

This paper provides an overview of the results of a continuous, 2-month experimental evaluation of all timing data provided by several GPS receivers. The primary purpose of this experiment was to provide measurement data facilitating fault modeling in our project SynUTC¹, which aims at external clock synchronization in fault-tolerant distributed real-time systems. As expected, the GPS receivers under test exhibited a wide variety of failures, ranging from transient omissions up to considerable deviations of the timing signals provided. Whereas those findings justify the appropriateness of our basic failure assumptions, it became nevertheless apparent that rerunning the experiment for a longer duration and with new brands/models of GPS receivers is advisable.

Keywords: experimental evaluation, GPS timing receivers, integrity, availability, reliability, faults, external clock synchronization, fault-tolerant distributed real-time systems.

1 Introduction

A rapidly increasing number of users rely upon GPS as the primary/sole means for acquiring highly accurate positioning and timing information, see [Dan97] for an overview. In fact, the introduction of the GPS caused a revolution in such different areas as navigation, land surveying, and time transfer, and provides enabling technology for many diverse applications in those areas. Despite of the unrivaled accuracy and reliability of GPS, however, applications eventually emerged that caused searching questions of what quality of service can actually be expected from the GPS.

Primarily driven by the stringent safety requirements of civil aviation [Dur90], integrity of the timing/positioning data provided by GPS receivers became a primary issue in GPS research. Several receiver autonomous monitoring (RAIM, see e.g. [BSK89],

[Mic95]) and failure detection and isolation (FDI, see e.g. [DL95], [BD94]) schemes were developed, which try to identify/remove apparently erroneous data. In the meantime, such techniques are pretty standard and incorporated in low-cost GPS receivers [GKKT95] as well.

Although RAIM/FDI techniques considerably increase the reliability of the positioning/timing output of a GPS receiver, there are nevertheless inherent limitations. As a remedy, augmentations of the GPS have been proposed, which distribute information from external monitoring of the GPS satellites via dedicated *GPS Integrity Channels* (GIC); see e.g. [BC94], [VMLE94]. Related standardization efforts like the *Wide Area Integrity Broadcast* (WIB) [DE94]; however, primarily target FAA's *GPS Wide-Area Augmentation System* (WAAS) and are, hence, not likely to be incorporated in commercial (low-cost) GPS receivers.

Whether RAIM/FDI/GIC is employed or not, however, there is the problem of assessing, i.e., determining, GPS integrity. In fact, any appropriate experimental evaluation is difficult since GPS failures are rare events —although manufacturers of GPS receivers are well aware of their occurrence, see [Dan97], [GKKT95]— that are not easy to observe and to trace back [Gol90] in practice. Comprehensive fault modelling is also difficult due to the fact that GPS is a complex system, with many potential sources of errors: Faults leading to an erroneous positioning/timing output of a GPS receiver may occur in the quite advanced receiver electronics [Die95], as well as in the GPS space and/or control segment. Note that it is even non-trivial to characterize the “normal” (fault-free) operation of GPS; see [Con93].

Most existing attempts to assess GPS integrity hence consider faults in the space segment (“satellite outages”) only. Relying on an *a posteriori* analysis of observational data, fault models like the ones of [DC90], [DC91], [PPP94] assume satellite failure probabilities of about 10^{-4} per hour. Simulation is the primary evaluation tool in the vast majority of related work [DL95], [BSK89], [GKKT95], [PPP94], [VMLE94]; we only know of a few papers [BD94], [Dyk92], [DBS93], [NT93] that deal with experimental evaluation as well.

Given those limitations, it is not too surprising that usual assessments of integrity are not particularly

¹This research is part of our project SynUTC, which has been supported by the Austrian Science Foundation (FWF) under grant no. P10244-ÖMA and is now continued under the START programme Y41-MAT.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE DEC 1997		2. REPORT TYPE		3. DATES COVERED 00-00-1997 to 00-00-1997	
4. TITLE AND SUBTITLE Long-Term Evaluation of GPS Timing Receiver Failures				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Technische Universitat Wien, Department of Automation, Tkeitlstral3e 1, A-1040 Vienna, Austria,				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADA418244. 29th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, Long Beach, CA, 2-4 Dec 1997					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 16	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

meaningful for even more demanding GPS applications. For example, when GPS receivers are used as the sole means for establishing a common notion of time in fault-tolerant distributed real-time systems, not even a single erroneous timing datum should occur or escape detection. By contrast, typical integrity-dependent applications like civil aviation usually only require one "certified" positioning/timing datum within a certain time interval. Reasoning about (or, as in our case, arguing against²) such "naïve" GPS applications, however, requires knowledge of the erroneous behavior of *any* timing datum provided by a GPS receiver — something that cannot be inferred from existing integrity reports.

Therefore, when we started our work on external clock synchronization in fault-tolerant distributed real-time systems [Sch94], we decided to conduct some experiments of our own to justify our undertaking and to facilitate the task of fault modeling. A continuous experimental evaluation of six commercially available GPS timing receivers was eventually performed from December 12, 1995 to February 15, 1996, which is comprehensively documented in [Hoe96].

This paper provides an overview of the most important results obtained from the analysis of the 778 MB of sampled data. It is organized as follows: In Section 2, we provide some information on fault-tolerant distributed real-time systems that will further motivate our study. Section 3 contains an overview of our measurement setup, Section 4 elaborates on the generation of the reference time scale. The actual results of our experimental evaluation are presented in Section 5, along with a number of graphs and charts contained in the appendix. Some conclusions and directions of further work provided in Section 6 eventually complete our paper.

2 Fault-Tolerant Real-Time Applications

It is well-known that designing distributed real-time systems is considerably simplified when local clocks displaying a common (global) notion of time are available at all computing nodes. Temporally ordered events are in fact beneficial for a wide variety of tasks, ranging from relating sensor data gathered at different nodes up to fully-fledged distributed algorithms, see [Lis93] for some examples. A synchronization tightness in the ms-range is usually sufficient here, although there are applications like [Mar96] that call for a μ s-range precision. To achieve this goal, it is common practice in industry to equip each node of the distributed system with a modular GPS timing receiver.

This solution, however, is not feasible when stringent fault-tolerance requirements are to be met. For example, it was noted in [GKKT95] that a prototype

²We are of course aware of the fact that any experimental evaluation can only provide a "snapshot" of possible failures, which is not sufficient for *confirming* a certain failure assumption. However, the outcome of an experiment can *invalidate* one, thereby increasing the *verisimilitude* (truthlikeness, see [Pop89, Chap. 10]) of the "surviving" assumptions, and this is why experiments are nevertheless appropriate.

TDMA communications system at Motorola eventually broke down due to a certain GPS failure. Similarly, real-time systems architectures like the one of [HW97] that simply phase-lock an internal clock to the 1 pps (1 pulse-per-second) output of a GPS timing receiver are bound to assign incorrect timestamps — taken arbitrarily in a one-second interval — if there is just a single incorrect 1 pps pulse. Obviously, to assess the overall reliability of such systems, the erroneous behavior of all timing data provided by a GPS receiver must be known.

To satisfy increased reliability requirements, fault-tolerant external clock synchronization techniques have been developed; see [Sch97] for the current status of research. Among these is our *interval-based clock validation* approach introduced in [Sch94] and further developed in a number of papers [SS97], [Scho97], [SSHL97], [HSS97], etc., which solves the external clock synchronization problem for large-scale, fault-tolerant distributed real-time systems. It is based on the idea of verifying whether the highly accurate, but possibly faulty, "authoritative time" supplied, e.g. by a GPS receiver, is consistent with some less accurate but reliable "validation time" backed up by all the nodes' local (quartz) clocks. If so, the distinguished time is accepted; otherwise, it is discarded and the nodes rely on the validation time instead. In essence, our clock validation technique simultaneously increases the fault-tolerance degree and decreases the total number of GPS receivers required in the distributed system. Consequently, it does not suffer from the "forest" of GPS antennas that would be required for a large (LAN-based) distributed system with dedicated GPS receivers. Still, information on all possible failures of GPS receivers is mandatory here as well for proper fault modeling.

From the above considerations, it is apparent that one should know as much as possible about those (rare) events where a GPS timing receiver provides erroneous data. Getting this information, however, is difficult enough under "ideal" operating conditions, and in fact made worse by the requirement of covering issues like receiver (software) errors as well as suboptimal reception and interfacing conditions. After all, circumstances like a 250 μ s timing bias (as exhibited by our first NavSymm receiver) cannot be detected without external verification, which will usually never take place in real applications: Commonly, an off-the-shelf GPS receiver is connected to an interface, undergoes some configuring (without knowing the antenna's exact 3D-position) — and is then assumed to provide correct time.

Hence, we argue that a meaningful experimental evaluation need not care about "fairness" and "nominal test conditions" and similar things that are likely to be violated in practice. Consequently, we did not run our experiment with multiple instances of the same receiver, and did not bother too much with optimizing antenna positions (w.r.t., e.g., multipath problems) or detecting presumed signaling noise across the interface. We are convinced, though, that our results are more meaningful for practical purposes than those obtained

under “artificial” test conditions. Bear in mind, however, that our evaluation does not impose a legitimate ranking of the evaluated receivers’ *potential* capabilities.

3 Experimental Setup

Having decided on the general issues, our first task (Dec. 1994 – Jan. 1995) was selection and purchase of “representative” GPS receivers. Constrained by our limited budget, we were looking for an affordable collection of timing receivers that reasonably covers the available spectrum. Based on a market survey guided by [GPS94] and [GPS95], we eventually chose the models listed in Table 1.

The market survey was also used to decide the question of what kind of GPS interface our custom clock synchronization hardware, namely, the *Network Time Interface* (NTI) M-Module [HSS97] resp. its pivotal UTCSU-ASIC [SSHL97], should provide. Since it turned out that all timing receivers support

- a digital (usually TTL-level) 1 pps signal that accurately indicates the beginning of a second,
- a serial interface utilizing some (proprietary) protocol to supply the current *time tag* (minute, hour, day, year) as well as additional status information,

whereas a few high-end receivers also provide

- an additional digital *status signal* to indicate health of the 1 pps pulse,
- a 10 MHz *frequency output*,

we decided to support both the 1 pps and a status signal. Moreover, our NTI can optionally be paced with an externally supplied 10 MHz frequency.

The next step was to conceive questions that were to be answered by our experimental evaluation. With $x(t)$ denoting the deviation of a receiver’s view of time and a suitable reference time (ideally, t) at the occurrence time t of the appropriate 1 pps pulse, the most important ones can be stated as follows:

- (1) How does the distribution function of the deviations $x(t)$ look like?
- (2) How does the distribution function of $\Delta x(\tau) = x(t + \tau) - x(t)$ for some fixed τ , i.e., the difference between deviations lying τ seconds apart, look like?
- (3) Same as above for an “artificial” 1 pps pulse obtained by dividing the 10 MHz frequency output (if provided) by 10^7 .
- (4) Are there missing or faulty 1 pps pulses and how does $x(t)$ behave in such cases?
- (5) Is there wrong information (time tag, health status) provided via the serial interface?

We will justify the appropriateness of the above questions in Section 5, where we discuss our results.

For data acquisition, answering those questions implies collection of the one-second timing information of

all GPS receivers for the full period of measurement. This means that each 1 pps pulse (including the artificial ones derived from the 10 MHz outputs) of each GPS receiver must be “timestamped” according to a reference clock and stored for later processing. In addition, the information provided via the receivers’ serial interfaces must be correctly associated and saved with those timestamps.

In reality, this problem becomes difficult due to the fact that one has to do this simultaneously for several 1 pps signals and with a few ns resolution. Of course, there are manufacturers offering highly sophisticated equipment for measuring phase differences of a single input against a reference signal. Our limited budget, however, did not allow us to replicate such expensive equipment.

The eventually chosen experimental setup shown in Figure 1 thus incorporates the following standard components only:

- On top of the figure, there are the *GPS Receivers* under test (see Table 1), along with the $10^7 : 1$ prescalers for the 10 MHz outputs provided by the Stellar and the NavSymm receiver.
 - The 10 MHz output of the *Reference Clock* (Ball Efratom FRS-C rubidium clock in a climatic housing) is divided by 10^5 resp. 10^7 to generate a 100 pps resp. 1 pps reference signal. In addition, another prescaler 8 : 1 is responsible for providing a symmetric trigger signal with period 8 s used for alternation control (see below).
 - At the heart of our setup are two *Logic Analyzers* LA1 and LA2 (HP 16500B) used for sampling the 1 pps pulses with 8 ns resolution. Note that we utilized the analyzers’ event trigger mode, where a timestamp of the internal LA clock is sampled into memory upon occurrence of a pulse at any input channel.
- Both logic analyzers alternately perform data acquisition and memory transfer to the Measurement PC via a HP-IB interface. More specifically, the leading edge of the 8 s trigger signal mentioned above initiates sampling at LA1 and causes LA2 to be read out, whereas the falling edge triggers sampling at LA2 and read out of LA1. This way, continuous measurement is accomplished.
- Apart from being responsible for reading out the memory of the logic analyzers via the HP-IB interface, the *Measurement PC* also provides the serial interfaces required for getting time tag and status information from the GPS receivers. It associates the 1 pps data from the LA with the additional information and computes a full reference timestamp by combining the LA clock timestamps with the sampled 100 pps rubidium pulses. Complete records, including “spurious pulse” data, are eventually sent to the Server PC via a RS-232 interface.

- Finally, the *Server PC* is responsible for storing the data sent by the Measurement PC on disk, one separate file for each hour of measurement. It is connected to the campus network for easy access and backup purposes.

4 Reference Time

In the course of our experimental evaluation, all 1 pps pulses from the GPS receivers were timestamped according to a free-running *reference clock*. As outlined in the previous section, this is basically a rubidium atomic clock augmented by the logic analyzers' high-resolution internal clocks. In view of the long measurement period vs. the relatively low stability of the rubidium clock, however, we cannot simply use those reference clock readings as a reference time. We need to establish an "artificial" *reference time* instead, which approximates real-time³ as accurately as possible. For this purpose, we take advantage of the fact that a great deal of the stochastic fluctuations of the 1 pps pulses of a high-quality GPS receiver vary not too slowly, so that they can be averaged out even by our low performance reference clock.

More formally, with $T(\text{clock}, t)$ denoting the time value some specified *clock* displays at real time t , the offset $x(t)$ of the best of our GPS receivers (Stellar) from the reference clock, observed at any time t when a 1 pps pulse from the Stellar occurs, can be written as

$$x(t) \equiv T(\text{Stellar}, t) - T(\text{Refclock}, t). \quad (1)$$

The offset $x(t)$ is impaired by both systematic (deterministic) deviations and stochastic, zero-mean noise originating in either the Stellar's 1 pps pulses or the reference clock. Apart from constant time and frequency offsets (irrelevant for our purposes), the sampled data reveal a systematic frequency drift $D_0 \approx 5 \cdot 10^{-18} \text{ s}^{-1}$, which is, however, small enough to be ignored. To analyze the stochastic part, we employ the well-known device of power spectral analysis (see [Ste85] for an overview and e.g. [Car86] for a thorough introduction): By means of Fourier analysis techniques, the (*one-sided*) *power spectral density* $S_x(f)$ depicted in Figure 2 can be computed from the sampled data $x(t)$, which gives the "signal power" per unit frequency at the particular center frequency f . Therefore, the area under $S_x(f)$ in a range $f_1 \leq f \leq f_2$ gives the proportion of the total signal power of $x(t)$ caused by its frequency contributions lying in $[f_1, f_2]$.

We should add that $S_x(f)$ was actually computed from the power spectral density $S_y(f)$ of the first differences of $x(t)$, which satisfy a well-known relation. The required $S_y(f)$ was obtained by applying a Fast Fourier Transform to the entire, properly Hanning-windowed data sample. As elaborated in [WPI89], this produces

³We actually used GPST (the inherent system time of the GPS) instead of UTC as our notion of real-time. Apart from an integer number of leap seconds, GPST is almost equivalent to UTC, since it is steered to follow the MasterClock of the United States Naval Observatory UTC(USNO, MC) with high accuracy.

a basically unbiased estimate of the spectrum at "reasonable" frequencies, since the Hanning window removes the distorting effect of spectral leakage. However, it introduces gross errors at the lowest frequencies, which can be corrected by exploiting the well-known fact [All87] that there should be a dominating random walk behavior originating in the rubidium clock. Note that the particular noise level was determined by some additional measurements of our rubidium clock against a cesium normal, which were conducted for verification purposes.

Figure 2 reveals a number of interesting facts about the noise actually present in $x(t)$:

- For frequencies up to 10^{-2} Hz (period of 100 s), we are primarily concerned with white phase noise caused by the granularity of our reference clock.
- For lower frequencies down to approximately 10^{-5} Hz (period of 100000 s), white phase noise originating from the 1 pps output dominates. It is primarily caused by the GPS' Selective Availability (SA), a deliberate distortion of orbital and clock data of the satellites, which has been implemented to reduce the accuracy available to civilian GPS users. Taking into account that the Stellar averages over 5 different satellites, we find the level of noise in accordance with observation of others [Tho93].
- There are also two interesting peaks in the spectrum: The first one appears at a frequency corresponding to a period of one day and is probably caused by the daily variations of ionospheric and tropospheric conditions. The second peak at a period of half a day is presumably a consequence of the second harmonic of the first peak and the 12-hour periodicity of the GPS satellite constellation, which can produce such variations in case of imperfectly known antenna positions.
- At even lower frequencies (with a period longer than a day), the random walk frequency noise of the rubidium clock dominates. Note that it hides any slowly varying noise of the 1 pps pulses.

This information on the noise of $x(t)$ can be exploited to establish a less noisy reference time, which will be developed subsequently. Our approach rests upon a good approximation of the offsets x_{GPST} an ideal GPS receiver (continuously displaying perfect GPST) would provide w.r.t. our reference clock, at arbitrary real-times t :

$$x_{\text{GPST}}(t) \equiv \underbrace{T(\text{GPST}, t)}_{=t} - T(\text{Refclock}, t), \quad (2)$$

recall that we chose GPST as our measure of real-time, hence $T(\text{GPST}, t) = t$. Since the actually sought offsets $z_{\text{any}}(t)$ of any GPS receiver and GPST can be obtained

via

$$z_{\text{any}}(t) \equiv T(\text{any}, t) - \underbrace{T(\text{GPST}, t)}_{=t} = x_{\text{any}}(t) - x_{\text{GPST}}(t), \quad (3)$$

where of course

$$x_{\text{any}}(t) \equiv T(\text{any}, t) - T(\text{Refclock}, t), \quad (4)$$

we only need a good approximation of $x_{\text{GPST}}(t)$ to compute a good approximation of $z_{\text{any}}(t)$. Note carefully that (3) is totally independent of the reference clock.

To be able to use our knowledge of the noise properties of the 1 pps pulses of the Stellar, we define

$$z_S(t) \equiv T(\text{Stellar}, t) - \underbrace{T(\text{GPST}, t)}_{=t} = z_S^0(t) + \zeta_S(t), \quad (5)$$

taken at any occurrence time t of a 1 pps pulse. Herein, $\zeta_S(t)$ is meant to cover only known stochastic deviations (white phase noise from SA), whereas $z_S^0(t)$ contains the unknown rest, i.e., systematic deviations and slowly varying stochastic deviations (hidden by the rubidium clock's random walk noise).

Similarly, we split the time deviations

$$z_R(t) \equiv T(\text{rb clock}, t) - \underbrace{T(\text{GPST}, t)}_{=t} = z_R^0(t) + \zeta_R(t) \quad (6)$$

of the rubidium clock in a stochastic part $\zeta_R(t)$ and a systematic part $z_R^0(t)$, which can be described around a certain moment in time t_r by a quadratic polynomial

$$z_R^0(t) = z_0(t_r) + y_0(t_r) \cdot (t - t_r) + \frac{D_0(t_r)}{2} (t - t_r)^2. \quad (7)$$

With the foregoing definitions, the measured offset $x(t)$ can of course be rewritten as

$$x(t) = z_S(t) - z_R(t). \quad (8)$$

With $\overline{f(t)}$ denoting the result of applying any linear averaging operation—to be fully specified later—to $f(t)$, we hence obtain

$$\overline{x(t)} = \overline{z_S(t)} - \overline{z_R(t)}. \quad (9)$$

Combining this with $x_{\text{GPST}}(t) = -z_R(t)$, which is apparent from comparing (2) and (6), we find

$$\overline{x(t)} = x_{\text{GPST}}(t) + \overline{z_S(t)} + (z_R(t) - \overline{z_R(t)}). \quad (10)$$

If the second and third term on the right hand side of (10) are small, we have found a good approximation of GPST. To compare this with the situation without averaging, we observe e.g. from (1) and (2) that

$$x(t) = x_{\text{GPST}}(t) + z_S(t). \quad (11)$$

Averaging thus yields a reduction of $z_S(t) \rightarrow \overline{z_S(t)}$, but introduces new contributions from the reference clock

$z_R(t) - \overline{z_R(t)}$. Hence, our goal must be to choose an averaging operation that provides a suitable tradeoff. For that purpose, we introduce the linear averaging operation

$$\overline{x(t)} \equiv \int_{-\infty}^{\infty} h(\tau) x(t - \tau) d\tau. \quad (12)$$

along with two reasonable constraints on the weighting function $h(\tau)$, namely

$$\int_{-\infty}^{\infty} h(\tau) d\tau = 1, \quad \text{and} \quad h(\tau) = h(-\tau). \quad (13)$$

Splitting $z_S(t)$ resp. $z_R(t)$ into their corresponding systematic and stochastic parts according to (5) resp. (6), we can express (10) as

$$\overline{x(t)} = x_{\text{GPST}}(t) + \overline{z_S^0(t)} + \underbrace{\left(z_R^0(t) - \overline{z_R^0(t)} \right)}_{\text{systematic}} + \underbrace{\left(\zeta_S(t) + (\zeta_R(t) - \overline{\zeta_R(t)}) \right)}_{\text{stochastic}}. \quad (14)$$

Since $z_S^0(t)$ was assumed to consist of the systematic errors and slowly varying stochastic components, we can reasonably infer that $\overline{z_S^0(t)} = z_S^0(t)$, provided that the weighting function $h(\tau)$ is such that it puts (almost) zero weight to large τ . For determining $\overline{z_R^0(t)}$, we apply the averaging operation (12) to (7) to obtain

$$\overline{z_R^0(t)} = z_0(t_r) + y_0(t_r) \cdot (t - t_r) + \frac{D_0(t_r)}{2} \left((t - t_r)^2 + \int_{-\infty}^{\infty} h(\tau) \tau^2 d\tau \right). \quad (15)$$

Deriving this result, we used the fact that the required even symmetry of $h(\tau)$ implies $\int_{-\infty}^{\infty} \tau h(\tau) d\tau = 0$.

Subtracting (15) from (7), we arrive at

$$z_R^0(t) - \overline{z_R^0(t)} = -\frac{D_0(t_r)}{2} \int_{-\infty}^{\infty} h(\tau) \tau^2 d\tau. \quad (16)$$

Choosing the reference point $t_r = t$, (14) can eventually be written as

$$\overline{x(t)} = x_{\text{GPST}}(t) + z_S^0(t) - \frac{D_0(t)}{2} \int_{-\infty}^{\infty} h(\tau) \tau^2 d\tau + \overline{\zeta_S(t)} + (\zeta_R(t) - \overline{\zeta_R(t)}). \quad (17)$$

Note that the systematic frequency drift $D_0(t)$ was determined to be about $D_0 \approx 5 \cdot 10^{-18} \text{ s}^{-1}$, independently of t , which yields a negligible contribution in (17).

Having settled the systematic parts in (17), we now turn our attention to the stochastic ones abbreviated by $\zeta(t) = \zeta_S(t) + (\zeta_R(t) - \zeta_R(t))$. Our purpose is still to fix the weighting function $h(\tau)$ to be used in the averaging operation. This is done by extracting the individual power spectral densities $S_S(f)$ of $\zeta_S(t)$ and $S_R(f)$ of $\zeta_R(t)$ from Figure 2, and choosing a function $h(\tau)$ that minimizes the spectral density $S_\zeta(f)$ of $\zeta(t)$, i.e., the (known) stochastic noise present in (10). Note that once $h(\tau)$ is determined, the sought approximation $\bar{x}(t)$ of $x_{\text{GPST}}(t)$ can be computed numerically from the sampled data without difficulties.

The power spectral density of the sole reference clock was modelled by

$$S_R(f) = \frac{10^{-30} \text{ Hz}^3}{f^4} + \frac{10^{-22} \text{ Hz}}{f^2}, \quad (18)$$

where the first term reflects the —deliberately overestimated (at minimum 6 dB higher)— random walk of our rubidium clock, cf. the lowest frequencies in Figure 2. The second term accounts for the rubidium clock's white frequency noise, which was found in our measurements against the cesium normal at frequencies higher than 10^{-4} Hz.

To characterize the noise properties of the sole 1 pps pulses of the Stellar, we simply assume that the white phase noise apparent in Figure 2 extends to frequencies lower than $f = 10^{-5}$ Hz as well. Hence, we simply cut off the random walk part to arrive at $S_S(f)$ depicted in Figure 3 below. After all, the reference clock's random walk noise dominates at the lowest frequencies, so that we cannot extract further information out of $S_\zeta(f)$. Any overlooked noise remains in the term $z_S^0(t)$, cf. (5). Note that this is also true for the two peaks in Figure 2, which fully survive the averaging process.

With those preparations, we can eventually attack the problem of choosing $h(\tau)$: Since all contributions to $\zeta(t)$ are reasonably assumed to be zero-mean, stationary power signals and $\zeta_S(t)$ and $\zeta_R(t)$ are obviously statistically independent, the usual superposition and filtering properties [Car86, p. 166ff] for power spectral densities apply. The filter function involved in $\bar{\zeta}_S(t)$ is just $h(\tau)$, the one in $\zeta_R(t) - \bar{\zeta}_R(t)$ reads $\delta(\tau) - h(\tau)$ with $\delta(\tau)$ denoting Dirac's delta-function, so that immediately

$$S_\zeta(f) = S_S(f)H(f)^2 + S_R(f)(1 - H(f))^2,$$

where $H(f)$ is the Fourier transform of $h(\tau)$; note that, because of the requested even symmetry of $h(\tau)$, we have $H(f) = H^*(f)$, i.e., a real-valued function.

Differentiating equation (19) and equating it to zero, we obtain the optimal (minimizing) choice of $H(f)$ as

$$H(f) = \frac{S_R(f)}{S_R(f) + S_S(f)}. \quad (19)$$

Essentially, it says —what intuitively seems clear— that whenever $S_S(f)$ dominates over $S_R(f)$, the function $H(f)$ should be close to zero, and when the opposite is

true, close to one. Hence, a computationally feasible $h(\tau)$ should be chosen, whose power transfer function $H(f)^2$ changes rapidly from one to zero at a certain corner frequency. We decided to use a two-sided exponential function

$$h(\tau) = \frac{1}{2T} e^{-|\tau|/T} \leftrightarrow H(f) = \frac{1}{1 + (2\pi fT)^2}, \quad (20)$$

for this purpose. Figure 4 shows the power transfer function $H(f)^2$ for this weighting function, which decays as f^{-4} for f large enough.

The only remaining problem was to fix the parameter T , which was done by computing $S_\zeta(f)$ for various values of T and choosing the one with the lowest root-mean-square value of $\zeta(t)$ (which can be calculated simply by integrating the spectrum). The eventually chosen value was

$$T = 3450 \text{ s} \quad \text{resulting in} \quad \sigma_\zeta = 6.5 \text{ ns}. \quad (21)$$

Table 3 finally provides the results of applying this averaging function to the sampled data according to (10). For comparison, we provide the findings for the non-averaged case (11) as well.

In both cases, we are confronted with a term of unknown value $z_S^0(t)$ resp. $\bar{z}_S^0(t)$, which barely differ for the chosen value of T . As mentioned earlier, it covers both unknown systematic errors like imperfectly known antenna position and deterministic delays in the receiver, as well as slowly varying stochastic errors originating in SA and variations of the tropospheric/ionospheric delays. Unfortunately, there is no other way of resolving this uncertainty but to establish a more accurate access to GPST, e.g., by means of common-view techniques.

5 Results

In this section, we survey the results of the analysis of the 778 MB of sampled data in order to answer the questions of Section 3. Additional information and further details may be found in [Hoe96].

5.1 Accuracy-related Quantities

According to item (1) of the list of questions in Section 3, we have to consider the difference $x(t)$ of a receiver's view of GPS time and our⁴ reference time observed at the occurrence of the 1 pps pulse. The sought distribution of $x(t)$ is in fact the quantity of primary interest in most evaluation reports on GPS receivers, cf. [KMB94]. In our clock validation framework, it is primarily required for deciding what maximum time uncertainty must be granted for a correct GPS receiver.

⁴Refer to Section 4 for the expected accuracy of our reference time, that is, its deviation from real-time t . Note that the appropriateness of our method of computing the reference time is also confirmed by the fact that the distribution of $x(t)$ for the Motorola GPS receiver in Figure 7 is in accordance with the results obtained in [KMB94].

Figure 7 in the appendix shows the distribution of $x(t)$ for all of our GPS receivers. Only correct 1 pps pulses were considered here, and any known systematic bias (e.g. resulting from the antenna cable delay) was removed beforehand. Moreover, the respective mean value \bar{x} was subtracted in the plots to make direct comparison easier. Table 2 summarizes the characteristic values of those distributions, namely, mean \bar{x} , standard deviation (root mean square) σ_x , and minimum and maximum offset ϵ_- and ϵ_+ taken relatively to \bar{x} .

Next, we turn our attention to questions (2) and (3) in the list of Section 3, which are devoted to the frequency stability of a GPS receiver for short averaging times $\tau \in [10 \text{ s} \dots 100 \text{ s}]$. The sought distribution of the differences $\Delta x(\tau) = x(t + \tau) - x(t)$ between the time deviation of a GPS receiver's 1 pps pulses lying (integer) τ seconds apart is particularly important for rate synchronization purposes: Our clock validation approach targets 1 μs synchronization tightness [HSS97], which makes it inevitable to synchronize not only clock states but also clock rates. In [Scho97], a suitable rate synchronization algorithm was introduced that requires periodic initiation with period $P_R \in [10 \text{ s} \dots 100 \text{ s}]$. For good performance, however, the intervals between any two initiation events should be as regular as possible.

Figure 8 and 9 in the appendix show⁵ the distributions of $\Delta x(\tau)$, $\tau \in \{30 \text{ s}, 100 \text{ s}\}$, for all the 1 pps and 10 MHz-derived outputs of our GPS receivers. Similarly as before, only correct pulses were considered here. Table 4 summarizes the characteristic values of those (zero-mean) distributions, namely, standard deviation σ_x , maximum value $\epsilon(\tau) = \max |\Delta x(\tau)|$, and corresponding maximum mean frequency deviation (stability) $\nu(\tau) = \epsilon(\tau)/\tau$.

Relating the results for the 1 pps outputs vs. the "artificial" 1 pps pulses derived from 10 MHz frequency outputs answers the question whether rate synchronization could benefit from GPS receivers with frequency output. However, it turns out that the additional effort is not worthwhile, at least for GPS receivers like the Stellar or the NavSymm: The ordinary 1 pps signal does not provide a significantly worse behavior.

5.2 Faulty Behavior

The issues of primary interest for fault tolerance are of course items (4) and (5) in the list of questions in Section 3. Owing to the fact that the receivers of Table 1 performed quite differently in this respect, we briefly report on the observed failures of each receiver separately.

Stellar GPS 100A

This GPS receiver produced no failures except 12 "spurious" 1 pps pulses, which appeared —partly in bursts—

⁵Note that the results in Figure 8 resp. 9 and Table 4 are considerably spoiled by the quantization noise caused by the relatively coarse granularity (8 ns) of our measurement setup. The vertical "stripes" appearing in the distributions are a visible sign of this problem, which obviously becomes less serious when τ is increased.

arbitrarily in between regular ones. We suppose that interfacing problems (possibly caused by ground loops) are responsible for this problem.

NavSymm NTFR-S

The 1 pps pulses of the originally shipped NavSymm receiver exhibited a constant bias of 250 μs ahead of UTC, which went completely unnoticed up to our first measurement epoch. It turned out that this problem was the result of a known(!) firmware bug, which was fixed in a replacement unit eventually used for actual evaluation.

The collected 1 pps data from the NavSymm receiver revealed 8 pulse jumps similar to the one shown in Figure 5: Initially, the so-called ACC-TIME value in the SFM-Message (indicating the 2σ -accuracy) jumped to zero, which means unacceptable accuracy. A few seconds later, the time offset $x(t)$ increased to several milliseconds and persisted in that gross error for about 1 minute. In addition, a few 1 pps pulses were occasionally lost during that period as well. Even worse, ACC-TIME returned to normal values of about 200–300 ns shortly after the pulse jump, although the 1 pps pulses were still extremely offset from UTC. Eventually, $x(t)$ instantaneously decreased to about 30 μs , from where it slowly approached normal values (within another minute).

Moreover, the NavSymm receiver also produced 6 pulse ramps like the one shown in Figure 6: The whole phenomenon lasted more than 1 minute and started with a change of ACC-TIME to zero, after which the offset $x(t)$ gradually grew to magnitudes around 1 μs . While the pulse was still offlying, ACC-TIME resumed displaying normal values, which were eventually also reached by $x(t)$.

Note that both kinds of failure forced us to exclude certain 1 pps pulses when computing the statistics of the NavSymm receiver in Section 5.1. More specifically, we discarded all pulses that occurred within 150 s after ACC-TIME became zero.

There was another, unexpected failure in the RS232-supplied time tag accompanying 2–3 consecutive 1 pps pulses at the beginning of certain days: The indicated day was too small. For example, the receiver said 1995 Dec 14 00:00:00, while it should have said 1995 Dec 15 00:00:00.

Motorola VP-Oncore

There were no failures except of ten omitted 1 pps pulses, which were lost along with their time tags.

Trimble SVeeSix-CM2

There were no failures in the 1 pps pulses, but several incorrect time tags during leap second insertion: The Trimble receiver outputs the time tag via its RS232 interface both in UTC and GPST. It happened that a leap second insertion was announced for the end of 1995, which changed the difference GPST-UTC from 10 s to

11 s. The Trimble receiver inserted the leap second incorrectly on 1996 Jan 01 00:00:00 GPST (\approx 1995 Dec 31 23:59:50 UTC) instead of 1996 Jan 01 00:00:00 UTC. Therefore, the UTC second information was incorrect for ten seconds.

Magellan Brain

During 0.3 % of our evaluation period, the Magellan receiver was in the so-called *coasting-mode*, where too few satellites are tracked to obtain correct measurements. The linearly growing offset $x(t)$ indicates that the receiver derives its 1 pps pulse directly from its local clock in this case. Due to its low quality, this implies large time deviations at the end of long coasting periods (the longest one lasted 994 seconds!), which are instantaneously corrected upon the transition to normal operation.

Apart from coasting, the receiver also showed very irregular behavior during times of bad accuracy, which are indicated by the supplied TFOM value (time figure of merit, giving the expected time error of the 1 pps pulse). For about 10 % of the measurement period, this value indicated time errors greater than 1 μ s. The actual behavior of $x(t)$, however, was much worse: Apart from ramp errors, sudden jumps of 10 μ s and 1 ms occurred quite frequently. Note that such jumps were also observed for TFOM values less than 1 μ s. As in case of the NavSymm receiver, those frequent failures forced us to discard erroneous pulses from the statistical analysis in Section 5.1. More specifically, we discarded all pulses with TFOM > 1 μ s, as well as all remaining ones with offset $x(t)$ > 1 μ s.

In addition, one out of thousand of the time tags transmitted via RS232 were found to be incorrect (usually off by one second). This usually appears for whole "blocks" of consecutive seconds, primarily in conjunction with erroneous 1 pps pulses. The maximum observed block length was 993 s.

Rockwell Microtracker

The 1 pps pulse data of this receiver revealed a bad accuracy ($x(t)$ in the 10 μ s-range, cf. Figure 7), which made it impossible to separate normal behavior and failures.

6 Conclusions and Future Work

Our results, as limited as they are due to their snapshot-like nature, reveal a number of interesting facts about the issues touched in Section 2. First of all, our findings confirm that trusting blindly in all timing data provided by a GPS receiver is definitely inappropriate for fault-tolerant applications. Moreover, since failures like systematic bias cannot be locally detected, redundant verification information is mandatory — and this is exactly the basic assumption underlying our interval-based clock validation scheme.

In addition, the following facts deserve attention:

- Transient omissions of 1 pps pulses are relatively frequent.
- One cannot always rely upon the health status provided by a GPS receiver, in particular after a non-health situation.
- The time tag can be wrong, making some kind of agreement mandatory.

On the other hand, we do not have enough information on erroneous 1 pps pulses for quantitative modelling, in the sense that we are unable to give meaningful failure probabilities. We can only infer that the probability of any transient failure is about 10^{-6} , without significant correlations between different receivers. Our data reveal actually three different classes of transient failures, namely

- "obviously" erroneous (spurious) pulses,
- "step" pulses, characterized by a sudden deviation from the correct time,
- "ramp" pulses, which drift away from correct time gradually (and usually slowly).

Clock validation can eliminate the former two but not the latter one, which are hence those failures where we actually need more information. Note that our observations correspond nicely to findings in integrity research; recall Section 1, where one distinguishes *step* and *ramp errors*: It is well-known that the latter failures are more difficult to iron out by RAIM/FDI than the former ones; see, e.g., [Mic95], [GKKT95], [BSK89].

In order to get meaningful information on failure probabilities, it is inevitable to extend the 2-month period of experimental evaluation considerably. Rerunning our evaluation is also required for keeping track with the rapidly maturing GPS receiver technology, which renders our 1993–1995 receivers partly as out-of-date models. The improved reliability of state-of-the-art brands/models of GPS receivers, however, calls for an even much longer period of data acquisition in order to get hold of failures.

To support longer evaluation periods, several deficiencies and limitations of our experimental setup must be removed:

- The reliability of the measuring equipment — involving numerous components — needs improvement, since 42 of the more than 10^6 readouts of LA memory failed. Although this does of course not invalidate our results (the resulting probability of overlooking an erroneous behavior is only about 10^{-10}), something should be done about it.
- Noise induced by ground loops seems to be a major source of problems, which might even affect the interfaces to the GPS receivers.
- Size and lacking robustness of the experimental setup prohibit an easy change of location, which is desirable both for varying reception conditions and "calibration" purposes (see next item).
- It turned out that the stability of our (relatively low cost) rubidium atomic clock is not exciting

and should be somewhat improved. Incorporating common-view techniques or moving the whole experimental setup to an institution with a cesium atomic clock, preferably one that contributes to UTC, would of course be a more appealing alternative.

- Some countermeasures against long-lasting disruptions due to unnoticed power failures should also be considered.

An appropriately improved and extended experimental evaluation will be incorporated in the comprehensive evaluation of our interval-based clock validation architecture [HSS97]. It will utilize a completely different data acquisition system based on a custom timestamping unit in conjunction with robust industrial VME-bus components, which is currently under development.

Acknowledgements

We gratefully acknowledge the efforts of our "hosts" Dietmar Loy and Martin Horauer, who helped us to cope with several problems arising in the design and implementation of the experimental setup.

References

- [All87] D.W. Allan. *Time and Frequency (Time-Domain) Characterization, Estimation, and Predication of Precision Clocks and Oscillators*, IEEE Trans. on Ultrasonics, Ferroelectrics, and Frequency Control, UFFC-34(6), November 1987.
- [BC94] S. Basker, I. Casewell. *Time and Clock Issues in WADGPS*, Proc. 7th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS-94), Alexandria, 1994, vol. 2, p. 1509-1518.
- [BD94] P. Brown, F. van Diggelen. *Experiences with RAIM, altitude-aided RAIM, and FDI: risks and benefits*, Proc. 3rd International Conference on Differential Satellite Navigation Systems (DSNS 94), London, 1994, vol. 2, p. 1-7.
- [BSK89] J.M. Brown, H. Schwartz, D. Kolody. *RAIM-an implementation study*, Proc. 2nd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS-89), Colorado Springs, 1989, p. 379-388.
- [Car86] A.B. Carlson. *Communication Systems*, 3rd. ed., McGraw Hill, 1986.
- [Con93] R. Conley. *GPS performance: what is normal?*, Navigation, 40(3), 1993, p. 261-281.
- [Dan97] P.H. Dana. *Global Positioning System (GPS) Time Dissemination for Real-Time Applications*, Journal of Real-Time Systems 12(1), January 1997.
- [DBS93] F. van Diggelen, A. Brown, J. Spalding. *Test results of a new DGPS RAIM software package*, Proc. 49th Annual Meeting of the Institute of Navigation, Alexandria, 1993, p. 641-646.
- [DC90] J.-M. Durand, A. Caseau. *GPS availability II. Evaluation of state probabilities for 21 satellite and 24 satellite constellations*, Navigation, 37(3), 1993, p. 285-296.
- [DC91] J.-M. Durand, T.M. Carlier. *GPS Continuity: Initial Findings*, Proc. 4th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS-91), 1991, p. 971-980.
- [DE94] A.J. Van Dierendonck, P. Enge. *RTCA SC-159 Wide Area Integrity Broadcast/Wide Area Differential GPS Status*, Proc. IEEE Position Location and Navigation Symposium, Las Vegas, 1994, p. 733-738.
- [Die95] A.J. Van Dierendonck. *Understanding GPS Receiver Terminology: A Tutorial*, GPS World, January 1995.
- [DL95] Ren Da, Ching-Fang Lin. *Failure Detection and Isolation Structure for Global Positioning System Autonomous Integrity Monitoring*, Journal of Guidance, Control, and Dynamics, 18(2), March-April 1995, p. 291-297.
- [Dur90] J.-M. Durand. *GPS inadequacies: comparative study into solutions for civil aviation*, Journal of Navigation, 43(1), 1990, p. 8-17.
- [Dyk92] K.L. Van Dyke. *RAIM availability for supplemental GPS navigation*, Navigation, 39(4), 1992/93, p. 429-443.
- [GKKT95] G.J. Geier, T.M. King, H.L. Kennedy, R.D. Thomas. *Prediction of the time accuracy and integrity of GPS timing*, Proc. 49th IEEE International Frequency Control Symposium, San Francisco, 1995, p. 266-274.
- [Gol90] K.L. Gold. *Ghost busting. The return from the dead of PRN 8 causes pseudo-range transients*, Proc. 3rd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GPS-90), Colorado Springs, 1990, p. 569-575.
- [GPS94] *GPS World Receiver Survey*, GPS World, January 1994, p. 38-56.

- [GPS95] *GPS World Receiver Survey*, GPS World, January 1995, p. 46–56.
- [Hoe96] D. Höchtl, *Fehlverhalten von GPS-Satellitenempfängern im Hinblick auf ihren Einsatz im Forschungsprojekt SynUTC* (in German), Master Thesis Dept. of Automation, Technische Universität Wien, 1996.
- [HSS97] M. Horauer, U. Schmid, K. Schossmaier. *NTI: A Network Time Interface M-Module for High-Accuracy Clock Synchronization*, submitted, 1997.
- [HW97] W.A. Halang, M. Wannemacher. *High Accuracy Concurrent Event Processing in Hard Real-Time Systems*, J. Real-Time Systems, 12(1), 1997.
- [KMB94] M. King, M. Miranian, D. Busch. *Test Results and Analysis of a Low Cost Core GPS Receiver for Time Transfer Applications*, Proc. National Technical Meeting of the Institute of Navigation, January 1994.
- [Lis93] B. Liskov. *Practical uses of synchronized clocks in distributed systems*, Distributed Computing, 6, p. 211–219, 1993.
- [Mar96] K.E. Martin. *Powerful Connections: High-Energy Transmission with High-Precision GPS Time*, GPS World, March 1996, p. 20–36.
- [Mic95] W.R. Michalson. *Ensuring GPS Navigation Integrity using Receiver Autonomous Integrity Monitoring*, IEEE Aerospace and Electronic Systems Magazine, 10, Oct. 1995, p. 31–34.
- [NT93] D.A. Nethropp, B. Tanju. *Test and evaluation of embedded GPS systems*, Proc. National Technical Meeting of the Institute of Navigation, Alexandria, 1993, p. 295–304.
- [Pop89] K. R. Popper. *Conjectures and Refutations*. 5th ed., Routledge, London.
- [PPP94] S.P. Pullen, B.S. Pervan, B.W. Parkinson. *A new approach to GPS integrity monitoring using prior probability models and optimal threshold search*, Proc. IEEE Position Location and Navigation Symposium, Las Vegas, 1994, p. 739–746.
- [Sch94] U. Schmid. *Synchronized UTC for Distributed Real-Time Systems*, Proc. IFAC WRTIP'94, Lake Constance, Germany, p. 101–107.
- [Sch95] U. Schmid. *Synchronized Universal Time Coordinated for Distributed Real-Time Systems*, Control Engineering Practice 3(6), p. 877–884, 1995. (Reprinted from [Sch94]).
- [Sch97] U. Schmid (ed.). Special Issue J. Real-Time Systems on *The Challenge of Global Time in Large-Scale Distributed Real-Time Systems*, Journal of Real-Time Systems, 12(1–3), 1997.
- [Scho97] K. Schossmaier. *An Interval-Based Framework for Clock Rate Synchronization*, to appear in Proc. ACM Symp. on Principles of Distributed Computing (PODC'97), 1997.
- [SS97] U. Schmid, K. Schossmaier. *Interval-based Clock Synchronization*, J. Real-Time Systems, 12(2), 1997, p. 173–228.
- [SSHL97] K. Schossmaier, U. Schmid, M. Horauer, D. Loy. *Specification and Implementation of the Universal Time Coordinated Synchronization Unit (UTC SU)*, Journal of Real-Time Systems, 12(3), May 1997.
- [Ste85] S.R. Stein. *Frequency and Time — Their Measurement and Characterization*, in E.A. Gerber, A. Ballato (eds.): Precision Frequency Control, Vol. 2, Academic Press, New York, 1985.
- [Tho93] C. Thomas. *Real-Time Restitution of GPS time*, 7th European Frequency and Time Forum, Neuchâtel, March 1993.
- [VMLE94] V.G. Virball, W.R. Michalson, P.L. Levin, P.K. Enge. *A GPS integrity channel based fault detection and exclusion algorithm using maximum solution separation*, Proc. IEEE Position Location and Navigation Symposium, Las Vegas, 1994, p. 747–754.
- [WPI89] F.L. Walls, D.B. Percival, W.R. Ireland. *Biases and Variances of several FFT Spectral Estimators as a Function of Noise Type and Number of Samples*, Proceedings of the 43rd Annual Symposium of Frequency Control, 1989.

Model	#Chn.	Year/Vers.	US \$
Stellar 100 GPS Clock	5	1995	4,500,-
NavSymm NTFR-S	6	1995/2.9	3,500,-
Motorola VP Oncore	6	1994/6.x	300,-
Trimble SVEeSix-CM2	6	1994	750,-
Magellan Brain	5	1993	600,-
Rockwell Microtracker	5	1993	350,-

Table 1: GPS receivers used for evaluation

Signal	\bar{x}/ns	σ_x/ns	ϵ_-/ns	ϵ_+/ns
Stellar 1 pps	0	37	-188	191
NavSymm 1 pps ^a	244	46	-263	361
Motorola 1 pps	108	46	-246	253
Trimble 1 pps	310	309	-1308	1346
Magellan 1 pps	60	155	-1602	1677
Rockwell 1 pps ^b	7.933	2.380	-5434	4518

Table 2: Characteristic values for the 1 pps pulses vs. reference time

^aNote that those characteristic values—in particular, the mean \bar{x} —are not fully compatible with the corresponding ones of the other receivers due to the fact that the NavSymm cannot output GPST on its 1 pps signal, but only UTC.

^bThe 1 pps pulse provided by the Rockwell receiver is not phase-locked to GPST but rather free-running. An offset value supplied via RS323 must be used to compute a “software 1 pps pulse”.

	Contribution	Amount	Origin/Reason for chosen value
without	$z_S^0(t)$	unknown	Systematic and slowly varying 1 pps noise
average	$\zeta_S(t)$	$\sigma = 38.0 \text{ ns}$	SA and quantization noise
	$z_S^0(t)$	$\approx z_S^0(t)$	Averaged $z_S^0(t)$
with	$\zeta_S(t)$	$\sigma = 5.7 \text{ ns}$	Averaged $\zeta_S(t)$
average	$z_R^0(t) - z_R^0(t)$	$\approx 0 \text{ s}$	No significant drift in data
	$\zeta_R(t) - \zeta_R(t)$	$\sigma = 3.2 \text{ ns}$	Introduced noise of rubidium clock

Table 3: Error contributions for non-averaged and averaged 1 pps pulses of Stellar receiver

Receiver	τ	$\sigma_{\Delta x}(\tau)/\text{ns}$	$\epsilon(\tau)/\text{ns}$	$\nu(\tau)$
Stellar	30 s	10.9	80	$2.7 \cdot 10^{-9}$
Stellar	100 s	27.2	196	$2.0 \cdot 10^{-9}$
Stellar 10 MHz	30 s	10.2	84	$2.8 \cdot 10^{-9}$
Stellar 10 MHz	100 s	27.0	200	$2.0 \cdot 10^{-9}$
NavSymm	30 s	27.1	201	$6.7 \cdot 10^{-9}$
NavSymm	100 s	47.5	399	$4.0 \cdot 10^{-9}$
NavSymm 10 MHz	30 s	29.7	224	$7.5 \cdot 10^{-9}$
NavSymm 10 MHz	100 s	49.2	394	$3.9 \cdot 10^{-9}$
Motorola	30 s	44.8	252	$8.4 \cdot 10^{-9}$
Motorola	100 s	51.7	290	$2.9 \cdot 10^{-9}$
Trimble	30 s	383.5	1688	$5.6 \cdot 10^{-8}$
Trimble	100 s	413.6	1715	$1.7 \cdot 10^{-8}$
Magellan	30 s	70.6	1896	$6.3 \cdot 10^{-8}$
Magellan	100 s	139.1	1865	$1.9 \cdot 10^{-8}$
Rockwell	30 s	210.4	6685	$2.2 \cdot 10^{-7}$
Rockwell	100 s	251.6	6997	$7.0 \cdot 10^{-8}$

Table 4: Characteristical values for pulses lying τ seconds apart

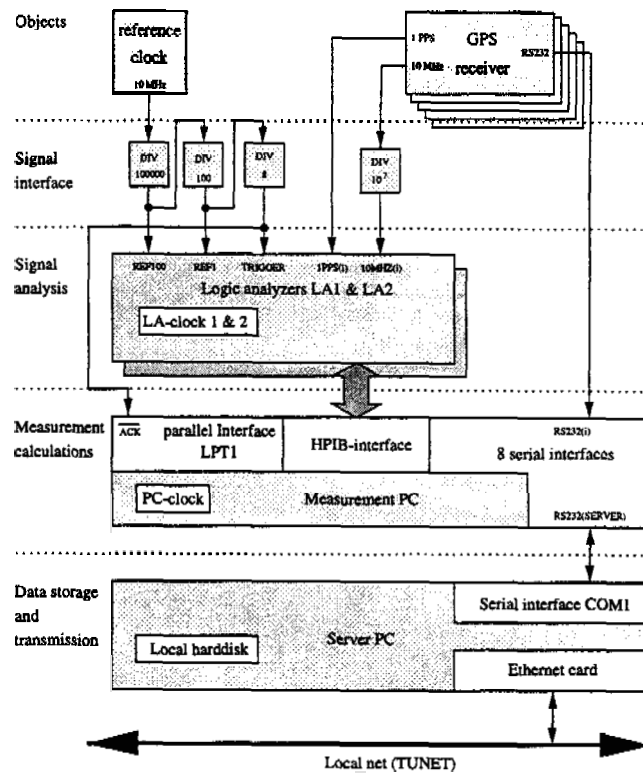


Figure 1: Schematics of experimental setup




Figure 2: Power spectral density for the Stellar's 1 pps pulses vs. reference clock




Figure 3: Power spectral density for 1 pps pulse cleaned from rb clock random walk noise




Figure 4: Power transfer function for a weighting function $h(\tau) = 1/2T e^{-|\tau|/T}$, $T = 3450$ s




Figure 5: Pulse jumps of the NavSym receiver




Figure 6: Pulse ramps of the NavSym receiver









Figure 7: Distribution functions for 1 pps pulses of all 6 receivers against the reference time (the means were eliminated)











Figure 8: Distribution functions for $x(t + \tau) - x(t)$ for all 6 receivers for $\tau = 30, 100$ s











Figure 9: Distribution functions for $x(t + \tau) - x(t)$ for all 6 receivers for $\tau = 30, 100$ s